

AN ADAPTIVE ALGEBRAIC MULTIGRID ALGORITHM FOR LOW-RANK CANONICAL TENSOR DECOMPOSITION*

HANS DE STERCK[†] AND KILLIAN MILLER[†]

Abstract. This paper presents a multigrid algorithm for the computation of the rank- R canonical decomposition of a tensor for low rank R . Standard alternating least squares (ALS) is used as the relaxation method. Transfer operators and coarse-level tensors are constructed in an adaptive setup phase based on multiplicative correction and on Bootstrap algebraic multigrid. An accurate solution is then computed by an additive solve phase based on the Full Approximation Scheme. Numerical tests show that for certain test problems the multilevel method significantly outperforms standalone ALS when a high level of accuracy is required.

Key words. multilevel method, algebraic multigrid, numerical multilinear algebra, canonical tensor decomposition, CANDECOMP, PARAFAC, alternating least squares

AMS subject classifications. 65N55 Multigrid methods, 65F10 Iterative methods, 15A69 Multilinear algebra

1. Introduction. In this paper we present a multigrid method for accurately computing a low-rank canonical decomposition of a tensor. An N th-order tensor is an N -dimensional array of size $I_1 \times \cdots \times I_N$ [20]. The *order* of a tensor is the number of modes (dimensions), and the size of the n th mode is I_n for $n = 1, \dots, N$. The canonical tensor decomposition is a higher-order generalization of the matrix singular value decomposition (SVD) in that it decomposes a tensor as a sum of rank-one components. For example, let \mathcal{T} be an arbitrary N th-order tensor of rank R , meaning that it can be expressed as a sum of no fewer than R rank-one components. Then its canonical decomposition is given by

$$\mathcal{T} = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \cdots \circ \mathbf{a}_r^{(N)}, \quad (1.1)$$

where \circ denotes the vector outer product. The r th rank-one component is formed by taking the vector outer product of N column vectors $\mathbf{a}_r^{(n)} \in \mathbb{R}^{I_n}$ for $n = 1, \dots, N$. For each mode $n = 1, \dots, N$, one can store the vectors $\mathbf{a}_1^{(n)}, \dots, \mathbf{a}_R^{(n)}$ as the columns of an $I_n \times R$ matrix $\mathbf{A}^{(n)}$. The matrix $\mathbf{A}^{(n)}$ is referred to as the *mode- n factor matrix* and its columns are the *mode- n factors*. The canonical tensor decomposition may then be expressed in terms of the factor matrices by the double bracket notation $[[\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}]]$, which is defined as the summation in (1.1).

We refer to the canonical decomposition as CANDECOMP/PARAFAC (CP) after the names originally given to it in early papers on the subject [12, 13]. Whereas (1.1) is an example of an exact CP decomposition, referred to as the *rank decomposition* since $R = \text{rank}(\mathcal{T})$, our goal is to find a CP decomposition for an arbitrary N th-order tensor \mathcal{Z} and a given number of components R . Depending on R the CP decomposition may only approximate \mathcal{Z} , in which case it is natural to look for the “best” such approximation in some sense. For example, this problem can be made

*This work was supported by NSERC of Canada and was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

[†]Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada (hdesterck@uwaterloo.ca, k7miller@uwaterloo.ca).

more concrete by posing it as an optimization problem:

$$\text{minimize } f(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) := \frac{1}{2} \left\| \mathcal{Z} - \llbracket \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)} \rrbracket \right\|^2. \quad (1.2)$$

In words, we seek the factor matrices that minimize the functional f . Here, $\|\cdot\|$ denotes the Frobenius norm of a tensor, defined as the square root of the sum of the squared entries of all tensor elements. In what follows, for matrices $\|\cdot\|$ refers to the Frobenius norm, and for vectors it refers to the vector two-norm. A general approach to solving the optimization problem (1.2) is to find a set of nontrivial factor matrices that zero out the gradient of f . In other words, at any local minimum of (1.2) the *first-order optimality equations* must be satisfied:

$$\frac{\partial f}{\partial \mathbf{A}^{(n)}} = \mathbf{0} \quad \text{for } n = 1, \dots, N. \quad (1.3)$$

In [1], Acar, Dunlavy and Kolda propose solving the first-order optimality equations by applying gradient-based optimization methods such as the nonlinear conjugate gradient method together with a line search. Along the same vein, a new method for computing the CP decomposition was recently developed by combining a nonlinear generalized minimal residual method with a line search [30]. In this paper we follow suit by solving the first-order optimality equations using a multigrid approach.

The optimization problem (1.2) is non-convex, consequently, it may admit multiple local minima. Moreover, for any local minimizer there is a continuous manifold of equivalent minimizers [20]. This manifold arises because of a scaling indeterminacy inherent to the CP decomposition, i.e., the individual factors composing each rank-one term can be rescaled without changing the rank-one term. The CP decomposition also exhibits a permutation indeterminacy in that the rank-one component tensors can be reordered arbitrarily [20]. In the following sections we discuss how the scaling and permutation indeterminacies can be removed by imposing a specific normalization and ordering of the factors. However, the CP decomposition may still exhibit multiple local minima for some problems, and depending on the initial guess, iterative methods may converge to different stationary points. Moreover, for certain tensors and certain values of R a best rank- R approximation does not exist [29]. Despite these difficulties, the exact CP decomposition has been shown to be unique up to scaling and permutation indeterminacies under mild conditions relating the ranks of the factor matrices with the tensor rank, and CP decompositions are used in many application fields [20].

The primary application of the CP decomposition is as a tool for data analysis, where it has been used in a variety of fields including chemometrics, data mining, image compression, neuroscience and telecommunications. A second class of problems is related to the decomposition of tensors arising from PDE discretization on high-dimensional regular lattices [27, 28, 17, 18]. Many algorithms have been proposed for computing the CP decomposition [20, 25, 22, 34, 1]. However, the workhorse algorithm today is still the original alternating least squares (ALS) method, which was first proposed in 1970 in early papers on the CP decomposition [12, 13]. The alternating least squares method is simple to implement, and often performs adequately, however, it can be very slow to converge, and its convergence may depend strongly on the initial guess. Despite its simplicity and potential drawbacks, it has proved difficult over the years to develop alternatives to ALS that significantly improve upon it in a robust way for large classes of problems. As a result, ALS-type algorithms are still considered the method of choice in practice.

The multigrid method described in this paper is intended to solve the first-order optimality equations for the CP optimization problem (1.3). It consists of two multilevel phases: a multiplicative correction scheme as the setup phase, and an additive correction scheme as the solve phase. In the setup phase a multiplicative correction scheme is used in conjunction with Bootstrap algebraic multigrid (BAMG) interpolation [6, 15] to not only build the necessary transfer and coarse-level operators, but also to compute initial approximations of the factor matrices. This phase uses the ALS method as the relaxation scheme on all levels. The setup phase is adaptive in the sense that the transfer operators are continually improved using the most recent approximation to the solution factor matrices. In order for the exact solution to be a fixed point of the multiplicative correction scheme, it needs to lie exactly in the range of interpolation at convergence. However, since each interpolation operator attempts to fit multiple factors (in a least-squares sense) this condition can be met only approximately. Therefore, after a few setup cycles we freeze the operators and use additive correction cycles in the solve phase, which can still converge when the exact solution lies only approximately in the range of interpolation. The combination of a multiplicative setup scheme and BAMG has already been considered in [21], where it formed the basis of an efficient eigensolver for multiclass spectral clustering problems. A similar approach was also proposed in [15, 7]. In the solve phase we use the Full Approximation Scheme (FAS) [5, 8, 11, 36] to efficiently obtain an accurate solution. Our multigrid framework is closely related to recent work on an adaptive algebraic multigrid solver for extremal singular triplets and eigenpairs of matrices [31], and to a lesser degree to multigrid methods for Markov chains [35, 32, 3]. Note that in this paper we give a complete but concise description of the proposed algorithm. Readers who are not familiar with adaptive algebraic multigrid (AMG) methods can find more background information about the general adaptive AMG approach in, for example, [9, 10, 31, 35].

The multigrid method proposed in this paper is expected to work well for tensors that have properties which make a multilevel approach beneficial, but not for generic tensors that lack these properties. Just as in the case of multigrid for matrix systems derived from PDE discretizations, our multilevel approach can lead to significant speedup when error components that are damped only weakly by the fine-level process can be represented and damped efficiently on coarser levels. We expect this to be the case for the decomposition of certain higher-order tensors that arise in the context of PDE discretization on high-dimensional regular lattices [27, 28, 17, 18], and we will illustrate the potential benefits of the proposed multigrid method for these types of problems. It should also be noted that the type of multigrid acceleration proposed in this paper will only be effective for low-rank decompositions with small R (e.g., up to 5 or 6). This is because a single interpolation operator is associated with an entire factor matrix, and each interpolation operator can only be expected to represent a small number of factors in a sufficiently accurate way, especially if the desired factors have little in common. These restrictions are entirely analogous to the case of the adaptive multigrid method for computing SVD triplets of a matrix [31].

The remainder of this paper is structured as follows. In §2 we define the basic notation and definitions used throughout this paper. Section 3 presents the first-order optimality equations and describes the alternating least squares method. Section 4 describes the multilevel setup phase, and §5 describes the multilevel solve phase. Implementation details and numerical results are presented in §6 followed by concluding remarks in §7.

2. Notation and definitions. This section outlines the notation that is used throughout this paper, much of which has been adopted from [1]. We also review some basic definitions and identities that are important to this paper. For further details we refer to the survey paper by Kolda and Bader [20], and the extensive references therein.

Vectors (tensors of order one) are denoted by boldface lowercase letters, e.g., \mathbf{v} . Matrices (tensors of order two) are denoted by boldface capital letters, e.g., \mathbf{A} . Higher-order tensors are denoted by boldface Euler script letters, e.g., \mathcal{Z} . The i th entry of a vector \mathbf{v} is denoted by v_i , element (i, j) of a matrix \mathbf{A} is denoted by a_{ij} , and, for example, element (i, j, k, ℓ) of a fourth-order tensor \mathcal{Z} is denoted by z_{ijkl} . The j th column of a matrix \mathbf{A} is denoted by \mathbf{a}_j . The n th element of a sequence is denoted by a superscript in parentheses, e.g., $\mathbf{A}^{(n)}$. In general, indices range from 1 to their capital versions, e.g., $n = 1, \dots, N$.

Matricization, also referred to as unfolding or flattening, is the process of reordering the elements of a tensor into a matrix. In this paper we are only interested in mode- n matricization, which arranges the mode- n fibers to be the columns of the resulting matrix. Note that a fiber is a higher-order analogue of matrix rows/columns, which is obtained by fixing every index of a tensor but one. Given a tensor \mathcal{Z} the mode- n matricized version is denoted by $\mathbf{Z}_{(n)}$.

Matricization provides an elegant way to describe the product of a tensor by a matrix in mode n . The n -mode matrix product of a tensor $\mathcal{Z} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ with a matrix $\mathbf{A} \in \mathbb{R}^{J \times I_n}$ is denoted by $\mathcal{Z} \times_n \mathbf{A}$ and is of size $I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N$. This product can be expressed in terms of unfolded tensors as follows

$$\mathcal{X} = \mathcal{Z} \times_n \mathbf{A} \quad \Leftrightarrow \quad \mathbf{X}_{(n)} = \mathbf{A} \mathbf{Z}_{(n)}.$$

For matrices $\mathbf{A} \in \mathbb{R}^{I \times K}$ and $\mathbf{B} \in \mathbb{R}^{J \times K}$, their *Khatri-Rao product* results in a matrix of size $(IJ) \times K$ defined by

$$\mathbf{A} \odot \mathbf{B} = [\mathbf{a}_1 \otimes \mathbf{b}_1 \dots \mathbf{a}_K \otimes \mathbf{b}_K],$$

where \otimes is the *Kronecker product*. These products have many useful properties, however, we will only need the associativity of the Khatri-Rao product, and the *mixed-product property* of the Kronecker product, i.e., $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$. We can then easily prove the following useful identity for any sequences of matrices $\mathbf{A}^{(n)}$ and $\mathbf{B}^{(n)}$, $n = 1, \dots, N$, of the appropriate sizes

$$\mathbf{A}^{(1)} \mathbf{B}^{(1)} \odot \dots \odot \mathbf{A}^{(N)} \mathbf{B}^{(N)} = \left(\mathbf{A}^{(1)} \otimes \dots \otimes \mathbf{A}^{(N)} \right) \left(\mathbf{B}^{(1)} \odot \dots \odot \mathbf{B}^{(N)} \right). \quad (2.1)$$

Another useful relationship between tensors and their matricized versions is as follows. Let $\mathcal{Z} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ and $\mathbf{A}^{(n)} \in \mathbb{R}^{J_n \times I_n}$ for $n = 1, \dots, N$. Then, for any $n \in \{1, \dots, N\}$

$$\mathcal{X} = \mathcal{Z} \times_1 \mathbf{A}^{(1)} \dots \times_N \mathbf{A}^{(N)} \quad \Leftrightarrow \quad \mathbf{X}_{(n)} = \mathbf{A}^{(n)} \mathbf{Z}_{(n)} \left(\mathbf{A}^{(N)} \otimes \dots \otimes \mathbf{A}^{(1)} \right)^T. \quad (2.2)$$

A proof of this property is given in [19].

To denote the product of a tensor and a sequence of matrices over some nonempty subset of the modes $\mathcal{N} = \{n_1, \dots, n_k\} \subset \{1, \dots, N\}$, we use the following notation as shorthand

$$\mathcal{Z} \times_{n \in \mathcal{N}} \mathbf{A}^{(n)} = \mathcal{Z} \times_{n_1} \mathbf{A}^{(n_1)} \dots \times_{n_k} \mathbf{A}^{(n_k)}.$$

3. CP first-order optimality equations and alternating least squares.

The first-order optimality equations for the CP decomposition are obtained by setting the gradient of the functional in (1.2) equal to zero. Following [1], for each mode $n \in \{1, \dots, N\}$ the derivative of f with respect to $\mathbf{A}^{(n)}$ can be written as an $I_n \times R$ matrix

$$\mathbf{G}^{(n)} = -\mathbf{Z}_{(n)}\Phi^{(n)} + \mathbf{A}^{(n)}\Gamma^{(n)}, \quad (3.1)$$

where

$$\Phi^{(n)} = \mathbf{A}^{(N)} \odot \dots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \dots \odot \mathbf{A}^{(1)}, \quad (3.2)$$

and

$$\Gamma^{(n)} = \Upsilon^{(1)} * \dots * \Upsilon^{(n-1)} * \Upsilon^{(n+1)} * \dots * \Upsilon^{(N)} \quad (3.3)$$

with $\Upsilon^{(n)} = \mathbf{A}^{(n)T} \mathbf{A}^{(n)}$ for $n = 1, \dots, N$. We note that since the $R \times R$ matrix $\Gamma^{(n)}$ is the Hadamard (element-wise) product “*” of symmetric positive semidefinite matrices, it too must be symmetric positive semidefinite [33]. Moreover, if each $\mathbf{A}^{(n)}$ has full rank then $\Gamma^{(n)}$ will be symmetric positive definite (SPD). The first-order optimality equations are then given by

$$\mathbf{G}^{(n)} = \mathbf{0}, \quad n = 1, \dots, N. \quad (3.4)$$

These equations offer a simple way to describe the ALS method for the CP decomposition. One iteration of ALS is equivalent to applying one iteration of block nonlinear Gauss–Seidel (BNGS) to the optimality equations (3.4). Iterating through the modes sequentially, at the n th step the factor matrices are fixed for all modes except n , and the resulting linear least-squares problem is solved for $\mathbf{A}^{(n)}$. In particular, $\Gamma^{(n)}$ and $\Phi^{(n)}$ are updated and $\mathbf{A}^{(n)} \leftarrow \mathbf{Z}_{(n)}\Phi^{(n)}(\Gamma^{(n)})^\dagger$, where $(\Gamma^{(n)})^\dagger$ is the Moore–Penrose pseudoinverse of $\Gamma^{(n)}$. Owing to the scaling indeterminacy inherent in the CP decomposition, it is possible that during ALS some factors may tend to infinity while others may compensate by tending to zero, such that the rank-one components remain bounded. This behavior can be avoided by using a normalization strategy. After each complete ALS iteration the factors of the r th component are normalized according to

$$\mathbf{a}_r^{(n)} \mapsto \lambda_r \left(\frac{\mathbf{a}_r^{(n)}}{\|\mathbf{a}_r^{(n)}\|} \right) \quad \text{for } n = 1, \dots, N, \quad \lambda_r = \left(\|\mathbf{a}_r^{(1)}\| \dots \|\mathbf{a}_r^{(N)}\| \right)^{1/N} \quad (3.5)$$

for $r = 1, \dots, R$. This normalization equilibrates the norms of the factors of each component, i.e., $\|\mathbf{a}_r^{(1)}\| = \dots = \|\mathbf{a}_r^{(N)}\|$. The ALS algorithm described here is used as the relaxation method and coarsest-level solver in the setup phase (see §4). We note that upon completion of the ALS iterations the rank-one terms are sorted in decreasing order of the normalization factors λ_r .

4. Multiplicative setup phase. In this section we describe the multilevel hierarchy constructed in the setup phase of our solver. We use two-level notation to describe the interaction of two grids at a time. Coarse grid quantities will usually be denoted by a subscript “ c ”, except in cases where a superscript “ c ” improves readability. Fine grid quantities and intergrid transfer operators have neither subscripts nor superscripts. The multiplicative setup phase described here is similar in concept to the setup phase of [31] for computing the SVD of a matrix, and more details about this general concept can be found in [31].

4.1. Derivation of coarse-level equations. The fine-level equations are given by the gradient equations stated in §3, i.e.,

$$\mathbf{Z}_{(n)} \left(\mathbf{A}^{(N)} \odot \dots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \dots \odot \mathbf{A}^{(1)} \right) = \mathbf{A}^{(n)} \mathbf{\Gamma}^{(n)} \quad (4.1)$$

for $n = 1, \dots, N$. Suppose there exist N full rank operators $\mathbf{P}^{(n)} \in \mathbb{R}^{I_n \times I_{n,c}}$, $I_{n,c} < I_n$, such that $\mathbf{A}^{(n)}$ lies approximately in the range of $\mathbf{P}^{(n)}$, i.e., for each n , $\mathbf{A}^{(n)} \approx \mathbf{P}^{(n)} \mathbf{A}_c^{(n)}$ for some coarse-level variables $\mathbf{A}_c^{(n)} \in \mathbb{R}^{I_{n,c} \times R}$. We note that since each factor matrix has R columns, it is unlikely that we can achieve equality. Then it follows that the solution to (4.1) can be approximated by solving a coarse-level problem

$$\begin{aligned} & \mathbf{P}^{(n)T} \mathbf{Z}_{(n)} \left(\mathbf{P}^{(N)} \mathbf{A}_c^{(N)} \odot \dots \odot \mathbf{P}^{(n+1)} \mathbf{A}_c^{(n+1)} \odot \mathbf{P}^{(n-1)} \mathbf{A}_c^{(n-1)} \odot \dots \odot \mathbf{P}^{(1)} \mathbf{A}_c^{(1)} \right) \\ &= \left(\mathbf{P}^{(n)T} \mathbf{P}^{(n)} \right) \mathbf{A}_c^{(n)} \mathbf{\Gamma}_c^{(n)} \quad \text{for } n = 1, \dots, N, \end{aligned} \quad (4.2)$$

followed by interpolation. Here, $\mathbf{\Gamma}_c^{(n)}$ is defined as in (3.3) with

$$\mathbf{\Upsilon}_c^{(n)} = \mathbf{A}_c^{(n)T} (\mathbf{P}^{(n)T} \mathbf{P}^{(n)}) \mathbf{A}_c^{(n)} \quad \text{for } n = 1, \dots, N.$$

By property (2.1) the left-hand side of (4.2) can be written as

$$\mathbf{P}^{(n)T} \mathbf{Z}_{(n)} \left(\mathbf{P}^{(N)} \otimes \dots \otimes \mathbf{P}^{(n+1)} \otimes \mathbf{P}^{(n-1)} \otimes \dots \otimes \mathbf{P}^{(1)} \right) \mathbf{\Phi}^{(n)} \quad \text{for } n = 1, \dots, N,$$

with $\mathbf{\Phi}^{(n)}$ given by (3.2). It then follows by property (2.2) that we can write the coarse-level tensor as

$$\mathbf{Z}^c = \mathbf{Z} \times_1 \mathbf{P}_1^T \times_2 \mathbf{P}_2^T \cdots \times_N \mathbf{P}_N^T, \quad (4.3)$$

which is essentially a higher dimensional analogue of the Galerkin coarse-level operator that is commonly used in multigrid for the matrix case. Defining $\mathbf{B}^{(n)} = \mathbf{P}^{(n)T} \mathbf{P}^{(n)}$ for each mode n we can write the coarse-level version of (4.1) as follows

$$\mathbf{Z}_{(n)}^c \left(\mathbf{A}_c^{(N)} \odot \dots \odot \mathbf{A}_c^{(n+1)} \odot \mathbf{A}_c^{(n-1)} \odot \dots \odot \mathbf{A}_c^{(1)} \right) = \mathbf{B}^{(n)} \mathbf{A}_c^{(n)} \mathbf{\Gamma}_c^{(n)}. \quad (4.4)$$

By the full rank assumption on the interpolation operators it follows that $\mathbf{B}^{(n)}$ is SPD, and hence we can compute its Cholesky factorization $\mathbf{B}^{(n)} = \mathbf{L}^{(n)} \mathbf{L}^{(n)T}$, where $\mathbf{L}^{(n)}$ is an $I_{n,c} \times I_{n,c}$ nonsingular lower triangular matrix. The Cholesky factors can be used to transform (4.4), whereby one obtains an equivalent set of equations that correspond to the first-order optimality equations of a coarse-level CP minimization problem. Making the change of variables $\hat{\mathbf{A}}_c^{(n)} = \mathbf{L}^{(n)T} \mathbf{A}_c^{(n)}$ for each mode n , and by appealing to property (2.1), it follows that (4.4) can be written as

$$\hat{\mathbf{Z}}_{(n)}^c \left(\hat{\mathbf{A}}_c^{(N)} \odot \dots \odot \hat{\mathbf{A}}_c^{(n+1)} \odot \hat{\mathbf{A}}_c^{(n-1)} \odot \dots \odot \hat{\mathbf{A}}_c^{(1)} \right) = \hat{\mathbf{A}}_c^{(n)} \hat{\mathbf{\Gamma}}_c^{(n)},$$

where

$$\hat{\mathbf{Z}}_{(n)}^c = \mathbf{L}^{(n)-1} \mathbf{Z}_{(n)}^c \left(\mathbf{L}^{(N)-1} \otimes \dots \otimes \mathbf{L}^{(n+1)-1} \otimes \mathbf{L}^{(n-1)-1} \otimes \dots \otimes \mathbf{L}^{(1)-1} \right)^T,$$

and $\hat{\mathbf{\Gamma}}_c^{(n)} = \mathbf{\Gamma}_c^{(n)}$ for $n = 1, \dots, N$. Moreover, by property (2.2) the transformed coarse-level tensor is given by

$$\begin{aligned}\hat{\mathbf{z}}^c &= \mathbf{z} \times_1 \left(\mathbf{P}^{(1)} \mathbf{L}^{(1)-T} \right)^T \cdots \times_N \left(\mathbf{P}^{(N)} \mathbf{L}^{(N)-T} \right)^T \\ &= \mathbf{z} \times_1 \hat{\mathbf{P}}^{(1)T} \cdots \times_N \hat{\mathbf{P}}^{(N)T},\end{aligned}\quad (4.5)$$

where $\hat{\mathbf{P}}^{(n)} = \mathbf{P}^{(n)} \mathbf{L}^{(n)-T}$ for $n = 1, \dots, N$. Therefore, the coarse-level equations are equivalent to the gradient equations of the following coarse-level functional:

$$\hat{f}_c(\hat{\mathbf{A}}_c^{(1)}, \dots, \hat{\mathbf{A}}_c^{(N)}) := \frac{1}{2} \left\| \hat{\mathbf{z}}^c - \llbracket \hat{\mathbf{A}}_c^{(1)}, \dots, \hat{\mathbf{A}}_c^{(N)} \rrbracket \right\|^2. \quad (4.6)$$

Therefore, the coarse-level equations can be solved by applying ALS to minimize the coarse-level functional \hat{f}_c . Initial approximations of the transformed coarse-level variables may be obtained by applying a restriction operator $\hat{\mathbf{R}}^{(n)}$ to the current fine-level approximations. In this paper we follow the standard approach of defining the restriction operators as the transposes of the interpolation operators, i.e.,

$$\hat{\mathbf{R}}^{(n)} = \hat{\mathbf{P}}^{(n)T} \quad \text{for } n = 1, \dots, N. \quad (4.7)$$

After solving the coarse-level equations, the coarse-grid-corrected fine-level approximations are obtained via prolongation:

$$\mathbf{A}_{\text{CGC}}^{(n)} = \hat{\mathbf{P}}^{(n)} \hat{\mathbf{A}}_c^{(n)} \quad \text{for } n = 1, \dots, N. \quad (4.8)$$

We note that if $\mathbf{P}^{(n)}$ contained $\mathbf{A}^{(n)}$ exactly in its range, and if the coarse-level equations were solved exactly, then the coarse-grid-corrected solution would equal the exact fine-level solution:

$$\mathbf{A}_{\text{CGC}}^{(n)} = \hat{\mathbf{P}}^{(n)} \hat{\mathbf{A}}_c^{(n)} = \mathbf{P}^{(n)} \mathbf{L}^{(n)-T} \mathbf{L}^{(n)T} \mathbf{A}_c^{(n)} = \mathbf{P}^{(n)} \mathbf{A}_c^{(n)} = \mathbf{A}^{(n)} \quad \text{for } n = 1, \dots, N.$$

However, since these conditions only hold approximately, we expect (4.8) to yield an improved but not exact approximation to the fine-level solution. In particular, since the approximation properties of the interpolation operators deteriorate as the number of components is increased, we expect our method to perform well for a relatively small number of components R .

4.2. Bootstrap AMG V-cycles. In this section we describe how we use Bootstrap AMG [6, 15] to find initial approximations of the desired factor matrices, and adaptively determine the interpolation operators that approximately fit the factor matrices. We follow the approach outlined in [21, 15, 31].

We begin by describing the initial BAMG V-cycle. On the finest level we randomly choose n_t *test blocks*, where each test block is a collection of N randomly generated *test factor matrices* (TFMs) $\mathbf{A}_t^{(1)}, \dots, \mathbf{A}_t^{(N)}$. The reason we must consider test blocks instead of simply adding more columns to the factor matrices is that rank-one components of the best rank- R CP tensor approximation must be found simultaneously [20]; contrary to the best rank- R matrix approximation, one cannot obtain the best rank- R CP approximation by truncating the best rank- Q approximation with $Q > R$. We also start with a collection of N randomly generated *boot factor matrices* (BFMs) $\mathbf{A}_b^{(1)}, \dots, \mathbf{A}_b^{(N)}$, which serve as our initial guess to the desired factor matrices. We

note that the subscripts “ t ” and “ b ” serve only to distinguish between the test and boot factors.

In the downward sweep of the first BAMG V-cycle we relax on each of the test blocks separately and also on the BFMs using the ALS algorithm described in §3. We then coarsen each mode on the finest level and determine the interpolation operators $\mathbf{P}^{(1)}, \dots, \mathbf{P}^{(N)}$. The n th interpolation operator $\mathbf{P}^{(n)}$ fits the factors in the n th TFMs across all test blocks in a least-squares sense, such that these factors lie approximately in the range of $\mathbf{P}^{(n)}$. We also build the coarse-level operator $\hat{\mathbf{Z}}^c$ and restrict the TFMs and BFMs to the first coarse level. This process is then repeated recursively until some coarsest level is reached, from which point on we relax only on the BFMs.

In the upward sweep of the first cycle, starting from the coarsest level, we recursively interpolate the BFMs up to the next finer level which gives the coarse-grid-corrected approximation on that level. We then relax the coarse-grid correction (CGC) using ALS. This process continues until the CGC on the finest level has been relaxed by ALS.

The initial BAMG V-cycle can be followed by several additional BAMG V-cycles. These cycles are the same as the initial cycle except for one key difference. In the downward sweep the n th interpolation operator $\mathbf{P}^{(n)}$ fits the factors in the n th TFMs across all test blocks as well as the factors in the n th BFM. Since the BFMs serve as our initial approximation for the additive phase of the algorithm, they must be well represented by interpolation if the additive solve phase is to converge.

4.3. Interpolation sparsity structure: coarsening. Construction of the interpolation operators proceeds in two phases. In the first phase, the sparsity structure of $\mathbf{P}^{(n)}$ is determined by selecting a subset of the fine-level indices $\Omega_n = \{1, \dots, I_n\}$. This subset, denoted by \mathcal{C}_n , is the set of coarse indices for mode n (it has cardinality $I_{n,c} < I_n$). Fine-level points that are not selected to the coarse level are represented by the set of fine indices $\mathcal{F}_n = \Omega_n \setminus \mathcal{C}_n$. For each point $i \in \mathcal{F}_n$ we define a set of coarse interpolatory points \mathcal{C}_n^i , which contains coarse points that i interpolates from. For convenience we assume that the points in \mathcal{C}_n^i are labeled by their coarse-level indices. Furthermore, for any fine-level point $i \in \mathcal{C}_n$ we let $\alpha(i)$ denote its coarse-level index. The interpolation operator $\mathbf{P}^{(n)}$ is defined by

$$p_{ij}^{(n)} = \begin{cases} w_{ij}^{(n)}, & i \in \mathcal{F}_n \text{ and } j \in \mathcal{C}_n^i \\ 1, & i \in \mathcal{C}_n \text{ and } j = \alpha(i) \\ 0, & \text{otherwise,} \end{cases}$$

where the $w_{ij}^{(n)}$ s are the interpolation weights for mode n . The interpolation weights are determined by a least-squares process described in §4.4. In this paper we use standard geometric coarsening for each mode, whereby \mathcal{C}_n consists of the odd numbered points in Ω_n , and \mathcal{F}_n consists of the even numbered points. It follows that $\alpha(i) = (i+1)/2$ in this case. For each $i \in \mathcal{F}_n$ we define $\mathcal{C}_n^i = \{\alpha(i-1), \alpha(i+1)\}$ (coarse-level labels) except possibly at the right endpoint. This coarsening works well when the modes have approximately the same size, however, for tensors in which the sizes of some modes vary widely, a more aggressive coarsening for the larger modes may be considered. In §4.5 we discuss a straightforward approach to coarsening tensors with varying mode sizes. While the simple coarsening procedure discussed here works well for the test problems considered in §6 (PDE problems on high-dimensional regular lattices), more general coarsening algorithms for other types of tensors would

be desirable. The development of such algorithms remains an interesting topic of future research.

4.4. Least-squares determination of interpolation weights. Suppose that mode n has been coarsened and that \mathcal{C}_n and \mathcal{F}_n are given. Further suppose that the factors in the n th test factor matrices across all test blocks are stored as the columns of the $I_n \times Rn_t$ matrix \mathbf{U}_t , and let $\mathbf{U}_b = \mathbf{A}_b^{(n)}$ for the boot factor matrices. Following the approaches of [5, 21, 15, 31] we use a least-squares (LS) process to determine the interpolation weights in the rows of $\mathbf{P}^{(n)}$ that correspond to points in \mathcal{F}_n . We want to fit the weights of $\mathbf{P}^{(n)}$ such that the vectors in \mathbf{U}_t and \mathbf{U}_b (except in the first cycle) lie approximately in the range of $\mathbf{P}^{(n)}$. Let the columns of $\mathbf{U}_f = [\mathbf{U}_t \mid \mathbf{U}_b]$ hold the $n_f = R(n_t + 1)$ vectors to be fitted. Let \mathbf{u}_k be the k th column of \mathbf{U}_f . Let $\mathbf{u}_{k,c}$ be the coarse-level version of \mathbf{u}_k obtained by injection, and let $(\mathbf{u}_{k,c})_j$ be its value in the coarse-level point j . Also, let u_{ik} be the value of \mathbf{u}_k in the fine-level point i . The interpolation weights of each row that corresponds to a point in \mathcal{F}_n may now be determined consecutively by independent least squares fits. Consider a fixed point $i \in \mathcal{F}_n$ with coarse interpolatory set \mathcal{C}_n^i . We solve the following least-squares problem to determine the unknown interpolation weights $w_{ij}^{(n)}$,

$$u_{ik} = \sum_{j \in \mathcal{C}_n^i} w_{ij}^{(n)} (\mathbf{u}_{k,c})_j \quad \text{for } k = 1, \dots, n_f. \quad (4.9)$$

We make (4.9) over-determined in all cases by choosing $n_t > M_s/R$, where M_s is the maximum interpolation stencil size for any i on any level, i.e., $|\mathcal{C}_n^i| \leq M_s$. Owing to the standard geometric coarsening of each mode $M_s = 2$, and so it is sufficient to use $n_t = 2$ for any number of components $R > 1$. For a rank-one decomposition we must take $n_t \geq 3$.

In practice it is necessary to formulate (4.9) as a weighted least-squares problem, where the weights should bias the fit toward the boot factors. We do so as follows. For a fixed n , we compute the gradient $\mathbf{G}^{(n)}$ according to (3.1) using the factor matrices in the current test block. The weights w_r for the mode- n factor vectors are then defined as

$$w_r = \frac{\|\mathbf{A}_t^{(n)}\|^2}{\|\mathbf{G}^{(n)}\|^2} \quad \text{for } r = 1, \dots, R. \quad (4.10)$$

Note that for a given mode n and test block all factor vectors have the same weights. The weights for all test blocks are stored in the vector \mathbf{w}_t of length Rn_t . The weights for the boot factors are computed similarly and are stored in the vector \mathbf{w}_b of length R . The full vector of weights $\mathbf{w} \in \mathbb{R}^{n_f}$ is obtained by “stacking” \mathbf{w}_t on top of \mathbf{w}_b . Equation (4.10) stems from the observation that $\mathbf{G}^{(n)}$ is a residual for the n th factor matrix. Therefore, since the BFM should converge much faster than the TFM, the gradient norm for the BFM should be much smaller, and hence the weights should be larger. We note that weights corresponding to a single factor matrix are chosen identical in (4.10) since we do not want preferential treatment given to different factor vectors, but rather to entire factor matrices. In our implementation the small weighted least-squares problems with diagonal weight matrix $\mathbf{W} = \text{diag}(\mathbf{w})$ are solved by a standard normal equations approach.

4.5. Pseudocode. A pseudocode description for a multiplicative setup phase V-cycle with ALS as the relaxation scheme and coarsest-level solver is given by Algorithm

1. For the unfamiliar reader we note that a multigrid V-cycle is obtained by performing one recursive solve on each level. If two recursive solves are performed on each level then a multigrid W-cycle is obtained. In this paper we only consider V-cycles for the setup phase.

Algorithm 1: V-cycle for setup phase of CP decomposition (CP-AMG-mult)

Input: tensor \mathbf{Z} , BFM s $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$, TFM s

Output: updated BFM s $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$, updated TFM s

1. Compute the set $\mathcal{J}_c = \{n : I_n > I_{n,coarsest}\}$
if $\mathcal{J}_c \neq \emptyset$ **then**
 2. Apply ν_1 relaxations to TFM s in each test block and to $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$
for $n \in \mathcal{J}_c$ **do**
 3. Build the interpolation operator $\mathbf{P}^{(n)}$ (on first cycle only use TFM s)
 4. Let $\mathbf{B}^{(n)} \leftarrow \mathbf{P}^{(n)T} \mathbf{P}^{(n)}$
 5. Compute the Cholesky factorization $\mathbf{B}^{(n)} = \mathbf{L}^{(n)} \mathbf{L}^{(n)T}$
 6. Let $\hat{\mathbf{P}}^{(n)} \leftarrow \mathbf{P}^{(n)} \mathbf{L}^{(n)-T}$ and $\hat{\mathbf{R}}^{(n)} \leftarrow \hat{\mathbf{P}}^{(n)T}$**end**
 7. Compute the coarse BFM s and coarse TFM s according to (4.11)
 8. Compute the coarse-level tensor $\hat{\mathbf{Z}}^c \leftarrow \mathbf{Z} \times_{n \in \mathcal{J}_c} \hat{\mathbf{R}}^{(n)}$
 9. Recursive solve:
$$\{\hat{\mathbf{A}}_c^{(1)}, \dots, \hat{\mathbf{A}}_c^{(N)}\} \leftarrow \text{CP-AMG-mult}(\hat{\mathbf{Z}}^c, \tilde{\mathbf{A}}_c^{(1)}, \dots, \tilde{\mathbf{A}}_c^{(N)}, \text{coarse TFM s})$$
 10. Compute the CGC $\mathbf{A}^{(n)}$ for $n = 1, \dots, N$ according to (4.12)
 11. Apply ν_2 relaxations to $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$
else
 12. Apply ν_c relaxations to $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$**end**
-

The CP-AMG-mult algorithm recursively coarsens each mode until it reaches some predefined coarsest level. However, since the size of each mode I_1, \dots, I_N may differ, it is possible that some modes may reach their coarsest level sooner than others. We address this issue as follows. For each mode n we define a threshold $I_{n,coarsest}$ to be the maximum size of that mode's coarsest level. For any mode n such that $I_n > I_{n,coarsest}$ we continue coarsening, otherwise we do not coarsen that mode further. Let the modes that still require further coarsening be indexed by the set $\mathcal{J}_c = \{n : I_n > I_{n,coarsest}\}$, and let \mathcal{J}_c' denote its complement. Then at any given level for each $n \in \mathcal{J}_c'$ it follows that $\hat{\mathbf{P}}^{(n)} = \mathbf{I}^{(n)}$, where $\mathbf{I}^{(n)}$ is the $I_n \times I_n$ identity matrix. Setting $\hat{\mathbf{P}}^{(n)}$ equal to the identity for all $n \in \mathcal{J}_c'$ has the following implications. The coarse-level tensor can be computed by taking the product in (4.5) over the modes in \mathcal{J}_c , instead of for all $n = 1, \dots, N$. The coarse-level approximations of the BFM s are given by

$$\tilde{\mathbf{A}}_c^{(n)} = \begin{cases} \hat{\mathbf{R}}^{(n)} \mathbf{A}^{(n)}, & n \in \mathcal{J}_c \\ \mathbf{A}^{(n)}, & n \in \mathcal{J}_c' \end{cases} \quad \text{for } n = 1, \dots, N. \quad (4.11)$$

Similarly, the coarse-level approximations of the TFM s in each test block are computed by restricting only those factor matrices indexed by \mathcal{J}_c . Additionally, the coarse-

grid-corrected BFM are given by

$$\mathbf{A}_{\text{CGC}}^{(n)} = \begin{cases} \hat{\mathbf{P}}^{(n)} \hat{\mathbf{A}}_c^{(n)}, & n \in \mathcal{J}_c \\ \hat{\mathbf{A}}_c^{(n)}, & n \in \mathcal{J}'_c \end{cases} \quad \text{for } n = 1, \dots, N. \quad (4.12)$$

The size of the coarsest level plays an important role in how the multigrid method performs. If the coarsest level is too large then not enough work is done on the coarser levels and convergence will be slow. Conversely, choosing too small a coarsest level may negatively impact convergence, or in some cases may even cause divergence (as in [21, 15, 31]). In practice we find that choosing $I_{n, \text{coarsest}} \geq R$ for all n works well.

5. Full approximation scheme additive solve phase. The Full Approximation Scheme (FAS) [5] is the nonlinear analogue of the linear additive correction multigrid method. When applied to linear problems FAS reduces to the usual additive method, and so it is a more general multigrid solver. In this section we discuss how FAS can be used to obtain an additive correction method for the CP decomposition. Two-level notation is used to describe the interaction of two grids at a time with coarse-grid quantities denoted by a subscript “ c ”.

5.1. Coarse-level equations. Recall the finest-level equations (4.1), which can be expressed as

$$\mathbf{A}^{(n)} \mathbf{\Gamma}^{(n)} - \mathbf{Z}_{(n)} \left(\mathbf{A}^{(N)} \odot \dots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \dots \odot \mathbf{A}^{(1)} \right) = \mathbf{0}$$

for $n = 1, \dots, N$. Suppose we define nonlinear operators $\mathbf{H}^{(1)}, \dots, \mathbf{H}^{(N)}$ such that for any $n \in \{1, \dots, N\}$

$$\mathbf{H}^{(n)} : \mathbb{R}^{I_1 \times R} \times \dots \times \mathbb{R}^{I_N \times R} \rightarrow \mathbb{R}^{I_n \times R}$$

where

$$\mathbf{H}^{(n)}(\{\mathbf{A}\}) := \mathbf{A}^{(n)} \mathbf{\Gamma}^{(n)} - \mathbf{Z}_{(n)} \left(\mathbf{A}^{(N)} \odot \dots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \dots \odot \mathbf{A}^{(1)} \right). \quad (5.1)$$

Note that we use $\{\mathbf{A}\}$ as shorthand for $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$. Then the fine-level problem can be formulated as a system of nonlinear equations

$$\mathbf{H}(\{\mathbf{A}\}) := (\mathbf{H}^{(1)}(\{\mathbf{A}\}), \dots, \mathbf{H}^{(N)}(\{\mathbf{A}\})) = (\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(N)}), \quad (5.2)$$

where $\mathbf{F}^{(n)} = \mathbf{0}$ for $n = 1, \dots, N$ on the finest level. In order to apply the full approximation scheme we require a coarse version of (5.2). For each mode n we define the coarse operator

$$\mathbf{H}_c^{(n)}(\{\mathbf{A}_c\}) := \mathbf{A}_c^{(n)} \mathbf{\Gamma}_c^{(n)} - \hat{\mathbf{Z}}_{(n)}^c \left(\mathbf{A}_c^{(N)} \odot \dots \odot \mathbf{A}_c^{(n+1)} \odot \mathbf{A}_c^{(n-1)} \odot \dots \odot \mathbf{A}_c^{(1)} \right), \quad (5.3)$$

where $\hat{\mathbf{Z}}_c$ is the coarse-level tensor computed in the multiplicative setup phase. Then the coarse-level FAS equations are given by

$$\mathbf{H}_c(\{\mathbf{A}_c\}) := (\mathbf{H}_c^{(1)}(\{\mathbf{A}_c\}), \dots, \mathbf{H}_c^{(N)}(\{\mathbf{A}_c\})) = (\mathbf{F}_c^{(1)}, \dots, \mathbf{F}_c^{(N)}) \quad (5.4)$$

where

$$\mathbf{F}_c^{(n)} = \hat{\mathbf{R}}^{(n)}(\mathbf{F}^{(n)} - \mathbf{H}^{(n)}(\{\mathbf{A}\})) + \mathbf{H}_c^{(n)}(\{\tilde{\mathbf{A}}_c\}) \quad \text{for } n = 1, \dots, N, \quad (5.5)$$

and $\hat{\mathbf{R}}^{(n)}$ is the mode n restriction operator from the multiplicative setup phase. Here $\tilde{\mathbf{A}}_c^{(n)}$ is the coarse-level approximation of $\mathbf{A}^{(n)}$ obtained by restriction. Solving (5.4) for $\{\mathbf{A}_c\}$, the coarse-grid-corrected approximations on the fine level are given by

$$\mathbf{A}_{\text{CGC}}^{(n)} = \mathbf{A}^{(n)} + \hat{\mathbf{P}}^{(n)}(\mathbf{A}_c^{(n)} - \tilde{\mathbf{A}}_c^{(n)}) \quad \text{for } n = 1, \dots, N, \quad (5.6)$$

where $\hat{\mathbf{P}}^{(n)}$ is the interpolation operator from the multiplicative setup phase. Together, Equations (5.2) and (5.4) to (5.6) describe a FAS two-level coarse-grid correction scheme for the CP optimality equations. In the following sections we describe the relaxation scheme used in the solve phase, and give a pseudocode description of the multilevel CP-FAS algorithm.

5.2. Relaxation. We employ block nonlinear Gauss–Seidel (BNGS) as the relaxation scheme and coarsest-level solver for the CP-FAS algorithm (Algorithm 2). Applying BNGS to the equations in (5.2) is similar to applying ALS to the CP optimality equations. One iteration consists of iterating through the modes sequentially, where at the n th step $\mathbf{\Gamma}^{(n)}$ and $\mathbf{\Phi}^{(n)}$ are computed, and then $\mathbf{A}^{(n)}$ is updated by solving

$$\mathbf{A}^{(n)}\mathbf{\Gamma}^{(n)} = \mathbf{Z}_{(n)}\mathbf{\Phi}^{(n)} + \mathbf{F}^{(n)}. \quad (5.7)$$

When considering how to solve (5.7) for mode n , on any level, we recall the following fact: the exact solution is a fixed point of FAS if it is a fixed point of the relaxation scheme and the coarsest-level solver. Suppose we update $\mathbf{A}^{(n)}$ by post-multiplying the right-hand side of (5.7) by $(\mathbf{\Gamma}^{(n)})^\dagger$, which is a small $R \times R$ matrix. If $\mathbf{\Gamma}^{(n)}$ is nonsingular then its pseudoinverse is equivalent to its inverse, in which case there exists a unique solution and the fixed point is preserved. However, if $\mathbf{\Gamma}^{(n)}$ is singular then post-multiplying by the pseudoinverse will in general not preserve the fixed point. Therefore, we consider an alternative approach. We propose using a few iterations of Gauss–Seidel (GS) to update $\mathbf{A}^{(n)}$, which guarantees the fixed point property of our relaxation method. Moreover, a result by Keller [16] for positive semidefinite matrices states that if $\mathbf{\Gamma}^{(n)}$ has nonzero entries on its diagonal then GS must converge to a solution (there may be many) of (5.7). Owing to the structure of $\mathbf{\Gamma}^{(n)}$ this condition is equivalent to the fundamental requirement that the factor matrices have nonzero columns. Therefore, we can be confident that GS will converge regardless of whether or not $\mathbf{\Gamma}^{(n)}$ is singular. In practice we find that only a few GS iterations are necessary to obtain a sufficiently accurate solution to (5.7), and that further iterations do little to improve the relaxed approximation. In this paper we use ten GS iterations.

After each iteration of BNGS on the finest level the factor matrices are normalized according to (3.5). Due to the structure of the FAS equations, in particular the right-hand side in (5.2), the scaling indeterminacy is not present on the coarser levels and so normalizing there is unnecessary. We note that the permutation indeterminacy is also removed on the coarser levels because of the (nonzero) right-hand side. Therefore, the rank-one terms are sorted in decreasing order of the normalization factors λ_r only on the finest level.

5.3. CP-FAS algorithm. A pseudocode description for an additive solve phase V-cycle is given in Algorithm 2. While other cycling schemes such as W-cycles and F-cycles [36] are available to us, in this paper we use V-cycles. We assume that at any given level the current tensor, the index set \mathcal{J}_c , and the interpolation/restriction

operators from the setup phase are available to the algorithm. Note that the parameters (ν_1, ν_2, ν_c) may be different from those used during the setup phase.

Algorithm 2: V-cycle for solve phase of CP decomposition (CP-FAS)

Input: right-hand side matrices $\{\mathbf{F}\}$, factor matrices $\{\mathbf{A}\}$

Output: updated factor matrices $\{\mathbf{A}\}$

if *not on the coarsest level* **then**

1. Apply ν_1 relaxations to $\mathbf{H}(\{\mathbf{A}\}) = (\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(N)})$
2. Coarse initial guess:

$$\tilde{\mathbf{A}}_c^{(n)} \leftarrow \begin{cases} \hat{\mathbf{R}}^{(n)} \mathbf{A}^{(n)}, & n \in \mathcal{J}_c \\ \mathbf{A}^{(n)}, & n \in \mathcal{J}'_c \end{cases} \quad n = 1, \dots, N$$

3. Coarse right-hand side:

$$\mathbf{F}_c^{(n)} \leftarrow \begin{cases} \mathbf{H}_c^{(n)}(\{\tilde{\mathbf{A}}_c\}) + \hat{\mathbf{R}}^{(n)}(\mathbf{F}^{(n)} - \mathbf{H}^{(n)}(\{\mathbf{A}\})), & n \in \mathcal{J}_c \\ \mathbf{F}^{(n)}, & n \in \mathcal{J}'_c \end{cases} \quad n = 1, \dots, N$$

4. Recursive solve:

$$\{\mathbf{A}_c^{(1)}, \dots, \mathbf{A}_c^{(N)}\} \leftarrow \text{CP-FAS}(\{\mathbf{F}_c\}, \{\tilde{\mathbf{A}}_c\})$$

5. Coarse-grid correction:

$$\mathbf{A}^{(n)} \leftarrow \mathbf{A}^{(n)} + \begin{cases} \hat{\mathbf{P}}^{(n)}(\mathbf{A}_c^{(n)} - \tilde{\mathbf{A}}_c^{(n)}), & n \in \mathcal{J}_c \\ \mathbf{A}_c^{(n)} - \tilde{\mathbf{A}}_c^{(n)}, & n \in \mathcal{J}'_c \end{cases} \quad n = 1, \dots, N$$

6. Apply ν_2 relaxations to $\mathbf{H}(\{\mathbf{A}\}) = (\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(N)})$

else

7. Apply ν_c relaxations to $\mathbf{H}(\{\mathbf{A}\}) = (\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(N)})$

end

It is instructive to mention the differences of this additive solution phase and the additive phase in the SVD case [31]. For the SVD, singular vectors can be computed in separate V-cycles and FAS is not required since the singular values are updated in a top-level Ritz step. In the tensor case, all factor vectors need to be computed simultaneously in a single FAS V-cycle, and the weights λ_r from (3.5), which are in some sense equivalent to the singular values, are updated in these FAS cycles as well, making a Ritz step unnecessary.

5.4. Full multigrid FAS cycles. For some problems the initial guess provided by the multiplicative setup phase may be too far from the solution to yield a convergent additive solve phase. One way in which we can try to obtain a better initial guess to the fine-level problem is to use Full Multigrid (FMG) [4, 11, 36]. Full multigrid is based on *nested iterations* whereby coarse levels are used to obtain improved initial guesses for fine-level problems. At any given level the problem is first solved on the next coarser level after which the solution is interpolated to the current level to provide a good initial guess. This process naturally starts at the coarsest level and terminates at the finest. Once an initial guess to the finest-level problem has been

obtained we can apply repeated CP-FAS cycles to obtain an improved approximate solution. We use CP-FAS V-cycles as the solver on each level of the FMG cycle, except on the coarsest level where ALS is used (see §3). A pseudocode description of the FMG-CP-FAS algorithm is given in Algorithm 3. We assume that at any given level the current tensor, the index set \mathcal{I}_c , and the interpolation/restriction operators from the setup phase are available. In Algorithm 3 we use a subscript ℓ to index the current level, where $\ell = 0, \dots, L-1$. Note that a subscript ℓ on an interpolation operator indicates that level ℓ is mapped to level $\ell-1$.

Algorithm 3: Full multigrid cycle for solve phase of CP decomposition (FMG-CP-FAS)

Output: finest-level factor matrices $\mathbf{A}_0^{(1)}, \dots, \mathbf{A}_0^{(N)}$

1. On the coarsest level apply ν iterations of ALS with a random initial guess to obtain $\mathbf{A}_{L-1}^{(1)}, \dots, \mathbf{A}_{L-1}^{(N)}$
 2. Set $\ell \leftarrow L-1$
 - while** $\ell \neq 0$ **do**
 3. $\mathbf{A}_{\ell-1}^{(n)} \leftarrow \hat{\mathbf{P}}_{\ell}^{(n)} \mathbf{A}_{\ell}^{(n)}$ for $n = 1, \dots, N$
 4. $\{\mathbf{A}_{\ell-1}^{(1)}, \dots, \mathbf{A}_{\ell-1}^{(N)}\} \leftarrow \text{CP-FAS}(\{\mathbf{0}\}, \mathbf{A}_{\ell-1}^{(1)}, \dots, \mathbf{A}_{\ell-1}^{(N)})$
 5. $\ell \leftarrow \ell - 1$
 - end**
-

6. Implementation details and numerical results. In this section we present the results of numerical tests. All experiments are performed using MATLAB version 7.5.0.342 (R2007b) and version 2.4 of the Tensor Toolbox [2]. Timings are reported for a laptop running Windows XP, with a 2.50 GHz Intel Core 2 Duo processor and 4GB of RAM. Initial guesses for the boot factors and test factors are randomly generated from the standard uniform distribution. The initial boot factors are also used as the initial guess for the standalone ALS method. The stopping criterion for the numerical tests is based on the gradient of f . In particular, with

$$g(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) = \frac{1}{\|\mathbf{z}\|} \left(\sum_{n=1}^N \|\mathbf{G}^{(n)}\|^2 \right)^{1/2}, \quad (6.1)$$

where $\mathbf{G}^{(n)}$ is the mode- n partial derivative of f as defined in (3.1), we iterate until

$$g(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) < \tau, \quad (6.2)$$

or until the maximum number of iterations are reached. In this paper we perform at most 500 iterations of our multilevel method, and at most 10^4 iterations of ALS. The stopping tolerance is fixed at $\tau = 10^{-10}$. Table 6.1 lists the parameters used by the setup and solve phases. As in [21, 15, 31], a larger number of relaxations is required in the setup cycles to produce sufficiently accurate transfer operators.

For each numerical test we perform ten runs with a different random initial guess for each run. The values reported in the tables represent averages over the successful runs, where a run is deemed successful if the stopping criterion is satisfied prior to reaching the iteration limit. The tables compare the ALS method and the multilevel method with or without FMG-CP-FAS as part of the setup phase (see §6.1). For ALS we report the average number of iterations, the average execution time and the number of successful runs. For the multilevel method we report the average number

TABLE 6.1
CP-AMG-mult and CP-FAS parameters

Parameter	CP-AMG-mult	CP-FAS
Pre-relaxations ν_1	5	1
Post-relaxations ν_2	5	1
Relaxations on coarsest level ν_c	100	50
Cycle type	V-cycle	V-cycle
Number of test blocks n_t	2	n/a

of iterations (setup and solve phases), the average total execution time, the number of successful runs, the average speedup over ALS and the number of levels. The average speedup is determined as follows. For a given test and run, if both ALS and the multilevel method were successful, we divide the execution time of ALS by the execution time of the multilevel method to obtain the speedup for that run. These values are then averaged to obtain the average speedup for that test. We note that execution times do not include the evaluation of the stopping criterion.

6.1. Implementation details. The multilevel setup and solve phases have thus far been described separately, however, these phases can be combined in the following simple way. Since the factor matrices lie only approximately in the range of the interpolation operators, convergence of the setup cycles, as measured by the functional g , should stagnate after a few iterations. Therefore, after each setup cycle the current iterate $\{\mathbf{A}_{new}\}$ is compared to the previous iterate $\{\mathbf{A}_{old}\}$ and the setup cycles are halted once

$$g(\{\mathbf{A}_{new}\}) > (1 - \varepsilon)g(\{\mathbf{A}_{old}\}), \quad (6.3)$$

where the tolerance is set at $\varepsilon = 0.1$. Moreover, at most five setup cycles are performed, and (6.3) is checked only after three setup cycles have completed. Once the setup phase has completed, solution cycles are performed until the stopping criterion (6.2) is satisfied. To improve robustness, we can also try to detect stagnation of the solution cycles. After five solution cycles have completed, we check if $g(\{\mathbf{A}_{new}\}) \geq g(\{\mathbf{A}_{old}\})$ in each subsequent iteration. If this condition is satisfied then the current iterate $\{\mathbf{A}_{new}\}$ is discarded and the transfer operators are rebuilt by one down-sweep of CP-AMG-mult with the previous iterate $\{\mathbf{A}_{old}\}$ used for the boot factors. This process is carried out at most once, and any further indications of stagnation are ignored. Note that the boot factors are not updated by the down-sweep of CP-AMG-mult, as doing so would ruin any progress made by the solution cycles.

The combination of the setup and solve phases described above can be modified to include an FMG-CP-FAS cycle as part of the setup phase. After the setup cycles have completed, we perform one FMG-CP-FAS cycle to compute a new approximation to the boot factors. The transfer operators are then rebuilt using one down-sweep of CP-AMG-mult. Note that while the TFMs are updated by this process, the boot factors are not. We refer to this combination as ‘Multilevel + FMG’ in the tables and figures.

We conclude this section by considering the computational costs of one setup cycle, one solution cycle, and one FMG cycle. Let $\ell = 0, \dots, L - 1$ index the levels, and define

$$I_n^\ell := \frac{I_n}{2^\ell}, \quad P^\ell := \prod_{n=1}^N I_n^\ell = \frac{1}{2^{N\ell}} \prod_{n=1}^N I_n, \quad S^\ell := \sum_{n=1}^N I_n^\ell = \frac{1}{2^\ell} \sum_{n=1}^N I_n$$

for any $\ell \geq 0$. Assume for simplicity that \mathbf{Z} is dense and that each mode is coarsened at the same rate with L being the same for each mode. Consideration of Algorithm 1 shows that the most expensive operations on each level are the construction of the coarse-level operator, the relaxations, and the construction of the interpolation operators, in particular computing the weights for the least-squares fits. The coarse-level tensor is constructed by sequentially taking the n -mode product of the current tensor with the n th restriction operator for $n = 1, \dots, N$. Computing $\hat{\mathbf{Z}}^c$ on level ℓ requires $\mathcal{O}(P^\ell S^\ell)$ operations. The dominant computation for the relaxations and least-squares weights is the matrix product $\mathbf{Z}_{(n)} \Phi^{(n)}$. Since $\mathbf{Z}_{(n)}$ is of size $I_n^\ell \times (P^\ell/I_n^\ell)$ and $\Phi^{(n)}$ is of size $(P^\ell/I_n^\ell) \times R$ on the ℓ th level, forming this product requires $\mathcal{O}(NP^\ell R)$ operations. Therefore, by summing over all the levels, to leading order one setup cycle requires approximately

$$\left[\left(\frac{2^N}{2^N - 1} \right) (n_t(\nu_1 + 1) + \nu_1 + \nu_2 + 1) + \frac{\nu_c}{2^{N(L-1)}} \right] \cdot \mathcal{O}(NPR) + \left(\frac{2^N}{2^N - 1} \right) \cdot \mathcal{O}(PS) \quad (6.4)$$

operations, where $P = P^0$ and $S = S^0$. We note that PS scales only slightly worse than linear in P , and in particular $NPI_{min} \leq PS \leq NPI_{max}$ where I_{min} and I_{max} are the sizes of the smallest and largest modes, respectively. Consideration of Algorithm 2 shows that the most expensive operations on each level are the relaxations and the construction of the right-hand sides. By a similar analysis, to leading order one solution cycle requires approximately

$$\left[\left(\frac{2^N}{2^N - 1} \right) (\nu_1 + \nu_2 + 1 + 1/2^N) + \frac{\nu_c}{2^{N(L-1)}} \right] \cdot \mathcal{O}(NPR) \quad (6.5)$$

operations. Similarly, it follows that to leading order one FMG-CP-FAS cycle requires approximately

$$\left[\left(\frac{2^N}{2^N - 1} \right)^2 (\nu_1 + \nu_2 + 1 + 1/2^N) + \frac{\nu + (L-1)\nu_c}{2^{N(L-1)}} \right] \cdot \mathcal{O}(NPR) \quad (6.6)$$

operations. Note that in (6.6) ν is the number of ALS iterations performed on the coarsest level and (ν_1, ν_2, ν_c) are the CP-FAS parameters. In general a solution cycle is significantly cheaper than a setup cycle because of the extra work required by a setup cycle to relax on the TFMs, the typically larger number of relaxations performed on each level of a setup cycle, and the added work of constructing the coarse-level tensors (i.e., the $\mathcal{O}(PS)$ term). If \mathbf{Z} is sparse then further savings are possible on the finest level. In particular, to leading order the cost of one relaxation reduces to NR times the number of nonzero elements in \mathbf{Z} . In our current framework the coarse tensors will in general be dense; multiplication by the inverted Cholesky factors as in (4.5) destroys any sparsity. Therefore, it may be interesting to consider alternative formulations of the coarse-level equations, for example, working directly with equations of the form in (4.4); see also [31].

6.2. Sparse tensor test problem. The first test problem we consider is the standard finite difference Laplacian tensor on a uniform grid of size s^d in d dimensions. This test problem yields an N -mode sparse tensor \mathbf{Z} of size $s \times s \times \dots \times s$ with $N = 2d$.

We can efficiently construct \mathbf{Z} by reshaping the $s^d \times s^d$ matrix

$$\mathbf{Z} = \sum_{k=1}^d \mathbf{I}_{\ell(k)} \otimes \mathbf{D} \otimes \mathbf{I}_{r(k)},$$

where $\mathbf{I}_{r(k)}$ is the $s^{k-1} \times s^{k-1}$ identity matrix, $\mathbf{I}_{\ell(k)}$ is the $s^{d-k} \times s^{d-k}$ identity matrix, and \mathbf{D} is the $s \times s$ tridiagonal matrix with stencil $[-1, 2, -1]$. While this test problem is somewhat pedagogical in nature, it offers a good starting point to illustrate our method. The parameters for the various test tensors that we consider are given in Table 6.2.

TABLE 6.2
Parameters for sparse problem.

test	problem parameters
1	$N = 4, s = 20, R = 4$
2	$N = 4, s = 20, R = 5$
3	$N = 4, s = 20, R = 6$
4	$N = 4, s = 50, R = 2$
5	$N = 4, s = 50, R = 3$
6	$N = 4, s = 50, R = 4$
7	$N = 6, s = 20, R = 2$
8	$N = 6, s = 20, R = 3$
9	$N = 6, s = 20, R = 4$
10	$N = 8, s = 10, R = 2$
11	$N = 8, s = 10, R = 3$

TABLE 6.3
Sparse problem. Average number of iterations and time (in seconds) until the stopping criterion is satisfied with stopping tolerance 10^{-10} . Here ‘it’ is the number of iterations, ‘spd’ is the multilevel speedup compared to ALS, ‘ns’ is the number of successful runs, and ‘levs’ is the number of levels.

test	ALS			Multilevel				Multilevel + FMG				
	it	time	ns	it	time	spd	ns	it	time	spd	ns	levs
1	1897	24.0	10	37	7.6	3.2	10	36	8.0	3.1	10	2
2	3329	53.6	8	64	15.1	4.5	10	42	11.8	5.1	10	2
3	3587	70.3	9	67	17.5	4.0	9	32	10.6	6.8	10	2
4	5457	105.9	10	123	40.5	2.7	10	120	41.2	2.6	10	4
5	5508	150.3	4	182	64.9	2.3	9	99	40.9	3.8	9	4
6	6788	244.0	3	150	62.0	5.2	10	136	58.7	5.4	9	4
7	1619	187.0	10	48	111.4	1.7	10	52	128.3	1.5	10	3
8	3481	610.0	10	72	164.8	3.7	10	70	177.9	3.5	10	3
9	4085	939.5	10	76	209.6	4.5	10	78	228.9	4.2	10	3
10	634	229.5	10	48	170.8	1.3	10	56	203.5	1.1	10	3
11	1743	943.3	10	39	402.3	2.3	10	43	474.0	2.0	10	3

The results of the tests for the sparse test case are given in Table 6.3. The results show that our multilevel approach is anywhere from two to seven times faster than ALS for this test problem. For tests 1 to 6 (order 4 tensors), larger speedups are observed for the multilevel method with FMG. However, for tests 7 to 11 (order 6 and 8 tensors) larger speedups are observed for the multilevel method without FMG. For

higher-order tensors the setup phase of multilevel method with FMG is considerably more expensive than the setup phase of the multilevel method without FMG. The multilevel variants demonstrate similar robustness to varying initial guesses for this problem, however, in general we expect the multilevel method with FMG to be the most robust option. We also observe the trend that for each grouping of tests in Table 6.3, the speedup tends to increase as the number of components R increases.

Figures 6.1 and 6.2 illustrate the convergence history of ALS and the multilevel method for one run of tests 3 and 9 in Table 6.3, respectively. These plots are typical of the performance observed for this test problem. We note that the spike in the ‘Multilevel + FMG’ curves is due to the initial approximation to the solution computed by the single FMG-CP-FAS cycle performed after the setup cycles.

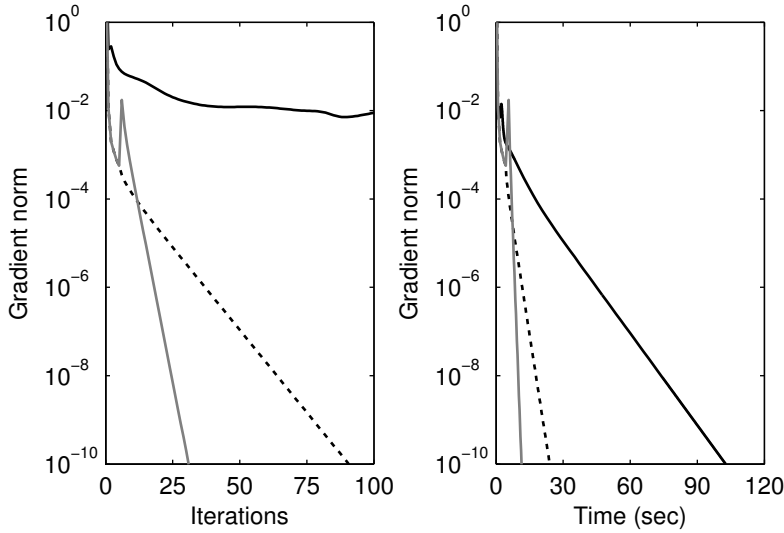


FIG. 6.1. *Sparse problem. Convergence plot for test 3 from Table 6.2 ($N = 4$, $s = 20$, $R = 6$). The solid black line is ALS, the solid gray line is the multilevel method with FMG, and the dashed line is the multilevel method without FMG.*

6.3. Dense tensor test problem. The second test problem we consider is a dense, symmetric third-order tensor $\mathbf{Z} \in \mathbb{R}^{s \times s \times s}$ whose elements are given by

$$z_{ijk} = (i^2 + j^2 + k^2)^{-1/2} \quad \text{for } i, j, k = 1, \dots, s.$$

This tensor was used as a test case in [27], which compares various methods for computing the CP decomposition including ALS. It was also considered in [28], which describes a novel method for computing the Tucker decomposition of third-order tensors. As mentioned in [28], \mathbf{Z} arises from the numerical approximation of an integral equation with kernel $1/\|\mathbf{x} - \mathbf{y}\|$ acting on the unit cube and discretized by the Nyström method on a uniform grid. In this section we compute CP decompositions for $R = 2, 3, 4, 5$. It has been observed numerically that when $R \geq 4$, ALS may be extremely slow to converge, for some initial guesses requiring on the order of 10^5 iterations, with highly non-monotonic convergence behavior. The performance of our method when $R \geq 4$ is less robust than desired because the multigrid framework uses a single interpolation operator for each factor matrix. Even so, depending on the initial guess our method may still demonstrate a significant improvement over ALS.

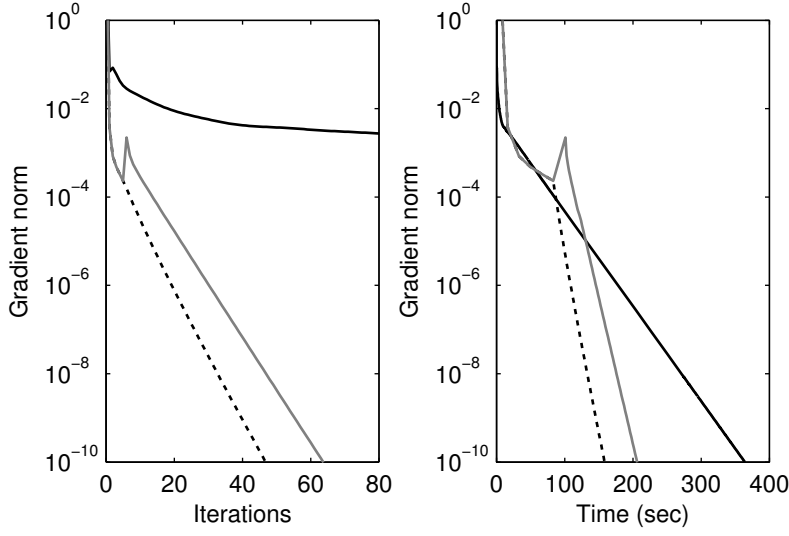


FIG. 6.2. *Sparse problem. Convergence plot for test 9 from Table 6.2 ($N = 6$, $s = 20$, $R = 4$). The solid black line is ALS, the solid gray line is the multilevel method with FMG, and the dashed line is the multilevel method without FMG.*

TABLE 6.4

Dense problem. Average number of iterations and time (in seconds) until the stopping criterion is satisfied with stopping tolerance 10^{-10} . Here ‘it’ is the number of iterations, ‘spd’ is the multilevel speedup compared to ALS, ‘ns’ is the number of successful runs, and ‘levs’ is the number of levels.

test	problem parameters	ALS			Multilevel + FMG				
		it	time	ns	it	time	spd	ns	levs
1	$s = 50$, $R = 2$	161	0.7	10	7	2.2	0.3	10	5
2	$s = 50$, $R = 3$	2435	11.8	10	15	2.7	4.7	10	5
3	$s = 50$, $R = 4$	4838	26.1	5	97	10.1	4.4	7	4
4	$s = 50$, $R = 5$				301	30.0		3	4
5	$s = 100$, $R = 2$	253	10.7	10	7	7.9	1.4	10	6
6	$s = 100$, $R = 3$	1695	80.2	9	9	8.1	10.1	10	6
7	$s = 100$, $R = 4$	3836	202.2	6	82	27.5	14.1	9	5
8	$s = 100$, $R = 5$	7854	455.2	2	192	62.0	9.0	4	5
9	$s = 200$, $R = 2$	274	90.3	10	7	48.8	1.9	10	7
10	$s = 200$, $R = 3$	1830	682.3	10	12	61.7	11.2	10	7
11	$s = 200$, $R = 4$	2998	1249.5	8	79	178.0	11.6	9	6
12	$s = 200$, $R = 5$	5686	2611.4	3	220	440.7	2.6	4	6

The results of the tests for the dense test case are given in Table 6.4. We note that only the multilevel method with FMG is considered for this test problem (see the description in §6.1). The blank entries for test 4 in Table 6.4 indicate that ALS did not have any successful runs. For $R \geq 3$ our multilevel approach can lead to significant savings in iterations and execution time. The speedup is less impressive when $R = 2$, since ALS already converges quickly without any multigrid acceleration. It is also apparent that as the number of components increases, the number of successful runs of the multilevel method, and of ALS, decreases. For initial guesses in which the multilevel method failed to converge, there was typically a rapid decrease in

the gradient norm, followed by convergence stagnation of the solution cycles. This behavior suggests that the setup phase was unable to construct transfer operators that adequately represented the solution in their range. Such cases were also characterized by slow convergence of ALS.

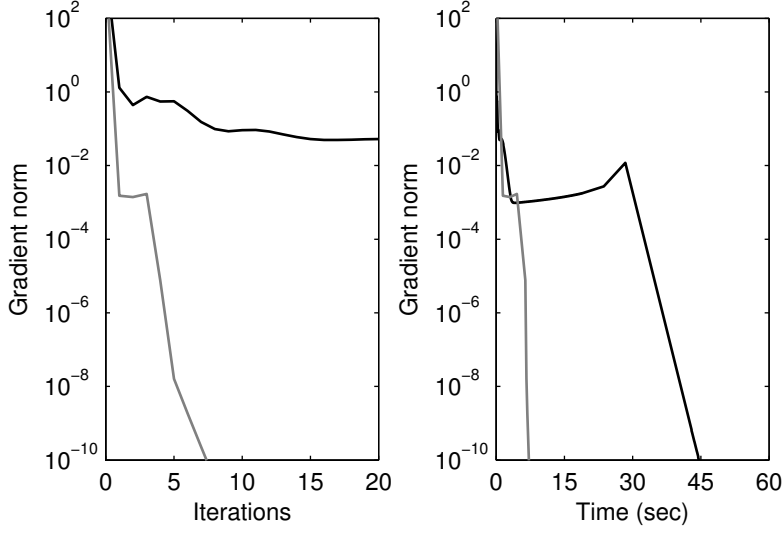


FIG. 6.3. Dense problem. Convergence plot for test 6 from Table 6.4 ($s = 100$, $R = 3$). The solid black line is ALS and the solid gray line is the multilevel method with FMG.

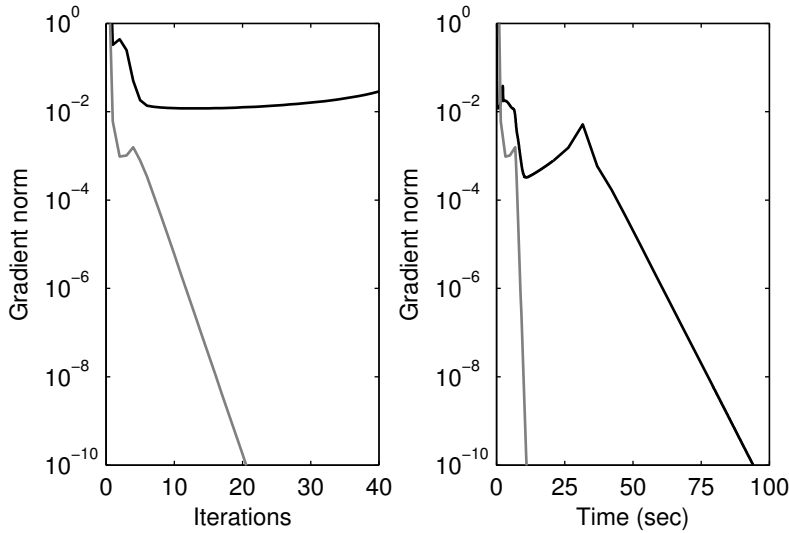


FIG. 6.4. Dense problem. Convergence plot for test 7 from Table 6.4 ($s = 100$, $R = 4$). The solid black line is ALS and the solid gray line is the multilevel method with FMG.

Figures 6.3, 6.4 and 6.5 illustrate the convergence history of ALS and the multilevel method for one run of tests 6, 7 and 8 in Table 6.4, respectively. Figure 6.5 ($R = 5$) shows how ALS can initially be slow to converge with erratic convergence

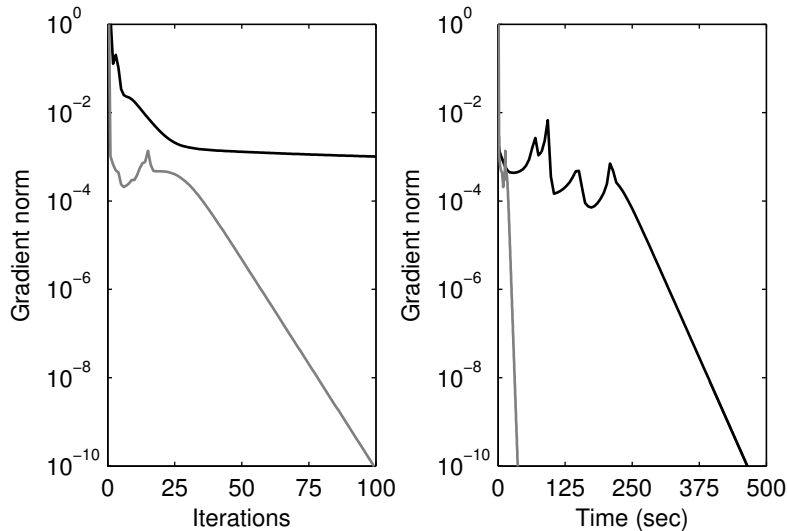


FIG. 6.5. *Dense problem. Convergence plot for test 8 from Table 6.4 ($s = 100$, $R = 5$). The solid black line is ALS and the solid gray line is the multilevel method with FMG.*

behavior: for the first half of the run its gradient norm fluctuates with little decrease. Such behavior can make it very difficult for the setup phase to construct adequate transfer operators.

7. Concluding remarks. We have presented a new algorithm for computing the rank- R canonical decomposition of a tensor for small R . As far as we are aware our method is the first genuine multigrid algorithm for computing the CP decomposition. Our work is also significant in that it presents the first adaptive AMG method for a nonlinear optimization problem. Similar to the method presented in [31] for computing SVD triplets of a matrix, we combined an adaptive multiplicative setup phase with an additive solve phase. The ALS method was used as the relaxation scheme. Numerical tests with dense and sparse tensors of varying sizes and orders (up to order 8) that are related to PDE problems showed how our multilevel method can lead to significant speedup over standalone ALS when high accuracy is desired.

Avenues of further research are plentiful. For example, it may be worthwhile to investigate a more sophisticated setup phase that iteratively combines the CP-AMG-mult cycles and CP-FAS-FMG cycles. As discussed briefly in §6.1, an alternative formulation of the coarse-level equations without the inverted Cholesky factors may be fruitful for sparse problems. In addition to PDE-related tensors, there may be other classes of tensors for which multigrid acceleration of ALS may be beneficial, but identifying and studying such classes remains a topic of future research. It would also be interesting to consider other ALS-type methods for the relaxations, for example, those using line searches, as well as to apply our method to the regularized optimization formulation of CP as described in [1]. Similarly, it would be interesting to generalize our multilevel framework to other similar tensor optimization problems such as the Tucker decomposition [20], block tensor decompositions [23, 26], best rank- (R_1, \dots, R_N) approximations [24, 14], and to other nonlinear optimization problems.

REFERENCES

- [1] E. ACAR, D. M. DUNLAVY, AND T. G. KOLDA, *A scalable optimization approach for fitting canonical tensor decompositions*, J. Chemometrics, 25 (2011), pp. 67–86.
- [2] B. W. BADER AND T. G. KOLDA, *Tensor toolbox for Matlab, version 2.4*. last accessed June, 2010.
- [3] M. BOLTEN, A. BRANDT, J. BRANNICK, A. FROMMER, K. KHAL, AND I. LIVSHITS, *A Bootstrap Algebraic Multigrid Method for Markov Chains*. manuscript, 2010.
- [4] A. BORZÌ AND G. BORZÌ, *Algebraic multigrid methods for solving generalized eigenvalue problems*, Int. J. Numer. Meth. Engng., 65 (2006), pp. 1186–1196.
- [5] A. BRANDT, *Multi-level adaptive solutions to boundary-value problems*, Math. Comput., 31 (1977), pp. 333–390.
- [6] ———, *Multiscale scientific computation: review 2000*, in Multiscale and Multiresolution Methods: Theory and Applications, Springer Verlag, Heidelberg, 2001, pp. 1–96.
- [7] A. BRANDT, J. BRANNICK, K. KAHL, AND I. LIVSHITS, *Bootstrap AMG*, SIAM J. Sci. Comput., 33 (2011), pp. 612–632.
- [8] A. BRANDT, S. F. MCCORMICK, AND J. RUGE, *Multilevel methods for differential eigenproblems*, SIAM J. Sci. Stat. Comput., 4 (1983), pp. 244–260.
- [9] M. BREZINA, R. D. FALGOUT, S. MACLACHLAN, T. A. MANTEUFFEL, S. F. MCCORMICK, AND J. RUGE, *Adaptive smoothed aggregation (α SA) multigrid*, SIAM J. Sci. Comput., 25 (2005), pp. 317–346.
- [10] ———, *Adaptive algebraic multigrid*, SIAM J. Sci. Comput., 27 (2006), pp. 1261–1286.
- [11] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, SIAM, Philadelphia, Pennsylvania, second ed., 2000.
- [12] J. D. CARROLL AND J. J. CHANG, *Analysis of individual differences in multidimensional scaling via an N-way generalization of “Eckart-Young” decomposition*, Psychometrika, 35 (1970), pp. 283–319.
- [13] R. A. HARSHMAN, *Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis*, UCLA working papers in phonetics, 16 (1970), pp. 1–84.
- [14] M. ISHTEVA, *Numerical methods for the best low multilinear rank approximation of higher-order tensors*, PhD thesis, K. U. Leuven, 2009.
- [15] K. KAHL, *Adaptive Algebraic Multigrid for Lattice QCD Computations*, PhD thesis, University of Wuppertal, 2009.
- [16] H. B. KELLER, *On the solution of singular and semidefinite linear systems by iteration*, J. Society for Industrial and Applied Mathematics: Series B, Numerical, 2 (1965), pp. 281–290.
- [17] B. KHOROMSKIJ, *On tensor approximation of Green iterations for Kohn-Sham equations*, Comput. Visual. Sci., 11 (2008), pp. 259–271.
- [18] ———, *Tensor-structured preconditioners and approximate inverse of elliptic operators in \mathbb{R}^d* , Constr. Approx., 30 (2009), pp. 599–620.
- [19] T. G. KOLDA, *Multilinear operators for higher-order decompositions*, Tech. Report SAND2006-2081, Sandia National Laboratories, Albuquerque, New Mexico and Livermore, California, 2006.
- [20] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.
- [21] D. KUSHNIR, M. GALUN, AND A. BRANDT, *Efficient multilevel eigensolvers with applications to data analysis tasks*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 32 (2010), pp. 1377–1391.
- [22] L. DE LATHAUWER, *A link between the canonical decomposition in multilinear algebra and simultaneous matrix diagonalization*, SIAM J. Matrix Anal. & Appl., 28 (2006), pp. 642–666.
- [23] ———, *Decompositions of a higher-order tensor in block terms – part II: Definitions and uniqueness*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 1033–1066.
- [24] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *On the best rank-1 and rank- (R_1, \dots, R_N) approximation of higher-order tensors*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1324–1342.
- [25] ———, *Computation of the canonical decomposition by means of a simultaneous generalized Schur decomposition*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 295–327.
- [26] L. DE LATHAUWER AND D. NION, *Decompositions of a higher-order tensor in block terms – part III: Alternating least squares algorithms*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 1067–1083.

- [27] I. V. OSELEDETS AND D. V. SAVOST'YANOV, *Minimization methods for approximating tensors and their comparison*, Comp. Math. Math. Phys., 46 (2006), pp. 1641–1650.
- [28] I. V. OSELEDETS, D. V. SAVOST'YANOV, AND E. E. TYRTYSHNIKOV, *Tucker dimensionality reduction of three-dimensional arrays in linear time*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 939–956.
- [29] V. DE SILVA AND LEK-HENG LIM, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM. J. Matrix Anal. & Appl., 30 (2008), pp. 1084–1127.
- [30] H. DE STERCK, *A nonlinear GMRES optimization algorithm for canonical tensor decomposition*, (2011). Submitted to SIAM J. Sci. Comput., arXiv: 1105.5331.
- [31] ———, *A self-learning algebraic multigrid method for extremal singular triplets and eigenpairs*, (2011). Submitted to SIAM J. Sci. Comput., arXiv: 1102.0919.
- [32] H. DE STERCK, K. MILLER, E. TREISTER, AND I. YAVNEH, *Fast multilevel methods for Markov chains*, Numer. Linear Algebra Appl., (2011). doi: 10.1002/nla.800.
- [33] G. STYAN, *Hadamard products and multivariate statistical analysis*, Linear Algebra Appl., 6 (1973), pp. 217–240.
- [34] G. TOMASI AND R. BRO, *A comparison of algorithms for fitting the PARAFAC model*, Comput. Stat. Data Anal., 50 (2006), pp. 1700–1734.
- [35] E. TREISTER AND I. YAVNEH, *On-the-fly adaptive smoothed aggregation multigrid for Markov chains*, SIAM J. Sci. Comput., 33 (2011), pp. 2927–2949.
- [36] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Elsevier Academic Press, San Diego, California, 2001.